

- 1 -

AUTOMATED ON-LINE CAPACITY EXPANSION METHOD FOR
STORAGE DEVICE

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to a method for successive capacity expansion during on-line operation
5 in a storage device system.

Description of the Related Art

Recently, storage domains of a plurality of host computers are often stored in a single storage device. This tendency is promoted by, for example, a
10 disk array. The disk array provides redundancy of the storage domains by a plurality of magnetic disk devices in a system, thereby increasing reliability and providing a necessary storage domain as a logical volume to a plurality of hosts. One of the merits of
15 using this disk array is capacity expansion of the logical volume.

For example, when a computer has used up a logical volume provided by a disk array, a free space in the disk array is cut out with an arbitrary length
20 to be allocated as a logical volume, so that the computer can connect this logical volume to a logical volume in use, thereby expanding the storage domain. This function of volume expansion can also be performed during on-line and is called on-line volume expansion.

The on-line volume expansion enables to expand a range of the logical volume storage domain corresponding to data increasing as time passes, without stopping the application and accordingly, enables to expand the application operation time. Moreover, upon volume capacity transfer, there is no need of data transfer between volumes, which significantly reduces the storage management cost.

Conventionally, a logical volume user should report to a logical volume provider when expanding the on-line volume. In a small-size site such as within a single enterprise, there is no sudden data increase and it is not necessary to perform on-line volume expansion all the time. However, in a large-size site such as a data center where a plurality of enterprises utilize data, there is a possibility of a sudden data increase from computers of the plurality of enterprises and it is necessary to perform the on-line volume expansion at all time. Moreover, in order to enhance the utilization effect of a storage domain of a disk array or the like, it is necessary to provide a volume to a plurality of users in a single storage device without waste.

In order to effectively use a storage domain in a storage device, it is necessary to manage the storage area in unit of a logical volume of a small capacity and to expand the logical volume with the small-capacity logical volume when required. When

using the on-line volume expansion in this condition,
there is a case that a request for the on-line volume
expansion is simultaneously caused by a plurality of
users due to the data increasing as time passes and the
5 logical volume provider may not be able to satisfy the
requests. In the worst case, the on-line volume
expansion cannot be performed and the computer
operation is stopped.

A logical volume user make a request for an
10 unlimited storage capacity without managing the storage
domain. On the other hand, the logical volume provider
should effectively manage the storage domain so as to
provide a logical volume to the user as rapidly as
possible without considering how the storage domain is
15 used by the logical volume user.

SUMMARY OF THE INVENTION

It is therefore an object of the present
invention to solve the requests of the storage resource
user and provider by using a storage domain management
20 apparatus capable of dynamically allocating a storage
domain of an appropriate capacity from the storage
domain collectively managed during a computer on-line
operation, thereby expanding the logical volume.

A computer system according to the present
25 invention includes at least one host computer, at least
one disk device, and a volume providing device which
are connected to one another. The volume providing

device managing a plurality of disk devices provides a logical volume corresponding to each of the host computers from the plurality of disk devices. The host computer sends an I/O request to a logical volume, and
5 a volume providing part interprets a logical block address where the I/O request is read and written. When the logical volume has no storage domain of the logical block address accessed by the I/O request, the volume providing device allocates a storage domain from
10 a free magnetic disk device, thereby dynamically expanding the logical volume storage domain. Moreover, the logical volume is reduced by an arbitrary amount indicated by an application.

BRIEF DESCRIPTION OF THE DRAWINGS

15 Fig. 1 shows configuration of a storage device system;

Fig. 2 shows a segment management table 2000;

Fig. 3 shows a physical and logical management table 3000;

20 Fig. 4 shows a physical relationship between LUN0 (4000), segment 1 (3410), segment 2 (3420), and segment 3 (3430);

Fig. 5 shows a flow of an I/O operation performed between a host 1100 and a volume server 1200;

25 Fig. 6 shows an operation flow of an I/O detector part 1231;

Fig. 7 shows an operation flow of a capacity

controller part 1232;

Fig. 8 shows an operation flow of a volume expansion process between a segment manager part 1233, the capacity controller part 1232, and physical and
5 logical address manager 1235;

Fig. 9 shows an operation of a volume reduction process between the segment manager part 1233, the capacity controller part 1232, and the physical and logical address manager 1235;

10 Fig. 10 shows an operation flow of segment state management by the segment manager part 1233;

Fig. 11 shows an operation flow of disk addition or deletion by the segment manager part 1233;

Fig. 12 shows an operation flow of an I/O
15 processor part 1234;

Fig. 13 shows an operation from of a physical and logical address conversion by the physical and logical address manager 1235;

Fig. 14 shows a logical relationship between
20 the LUN 14000, segment 14201, the segment 14202, and a magnetic disk device 14100;

Fig. 15 shows an operation flow of addition/deletion of a segment to/from the physical and logical management table by the physical and logical
25 address manager 1235;

Fig. 16 shows an operation flow of a logical volume storage capacity expansion between a commander part 1132 in the host 1100 and the capacity controller

part 1232 in the volume server 1200;

Fig. 17 shows a physical format of an I/O command; and

Fig. 18 shows an operation flow of a logical
5 volume storage capacity expansion between the commander
part 1132 in the host 1100 and the capacity controller
part 1232 in the volume server 1200.

DESCRIPTION OF THE EMBODIMENTS

Description will now be directed to
10 embodiments of the present invention with reference to
the attached drawings.

(1) First Embodiment

Fig. 1 shows a configuration m system
includes a plurality of hosts 1100, a volume server
15 1200, and a plurality of magnetic disk device 1300. In
this embodiment a single host 1100 and three magnetic
disk devices are used, but only one magnetic disc
device or any numbers of devices 1300 may be used
instead. Moreover, the host 1100, the volume server
20 1200, and the magnetic disk devices 1300 have CPU,
memory, and the like but they are not direction
relationship with explanation of the embodiment of the
present invention and their explanations are omitted.

The host 1100 has an application program
25 (App) 1110, an operating system (hereinafter, referred
to as OS) 1130, and a channel interface 1140. The
application 1110 is application software such as a DB

and a file system for reading and writing from/to a volume provided from the volume server 1200 as the volume providing device. The operating system 1130 includes a volume device part (rsd0) 1131 for receiving
5 an I/O request from the application 1110 and transferring the I/O request to the channel interface (I/F) 1140, and a commander part 1132. The commander part 1132 is software performing management of logical area information including a start block position and
10 an end block position of the storage device used by the OS 1130 and operation control of the volume server 1200.

By referring to Fig. 4 showing a logical volume as a storage domain provided by the volume
15 server 1200, explanation will be given on the logical area information 4100. Since in this embodiment a volume is dynamically expanded, the start block position of the logical area information handled by the OS 1130 is indicated by OLBA (4200) and the end block
20 position indicates a capacity of the maximum storage device supported by the OS of the storage device or a storage capacity used by a user. When the maximum capacity is limited by a storage capacity used by the user, the storage domain of the logical volume provided
25 by the volume server 1200 is also limited by this end block position. Moreover, the logical area information is stored in the storage device managed by the OS 1130 in a necessary area between the start block position

OLBA (4200) and LBA to define the necessary area (4100). Normal data to be stored is stored in a storage domain (4400) other than the logical area information.

5 The I/F 1140 shown in Fig. 1 is a block device I/F having a function to transfer an I/O from the OS 1130 to an external device such as a fibre channel I/F and an SCSI I/F.

 A command controlling between the host 1100
10 and the volume server 1200 and between the volume server 1200 and the magnetic disk device 1300 is called an I/O request. The I/O request is a command of a block device protocol such as the SCSI protocol. Fig. 17 shows a command format of the I/O request. The I/O
15 request has an operation code 17001, a LUN (logical unit number) 17002, an LBA (logical block address) 17003, and a transfer data length 17004. The operation code 17001 is a number indicating a read process, write process, and the like. The LUN 17002 is a unique
20 number of a logical volume to be processed by the command. The LBA 17003 is a position for processing a logical volume provided by the volume server. The transfer data length 17004 shows an amount to be processed all-at-once by the I/O request.

25 The volume server 1200 includes: a channel I/F 1210 at the host side; a volume provider part 1230 performing processing of an I/O request requested from the host 1100 and management of the magnetic disk

device 1300; and a channel I/F 1220 for transferring
the I/O request requested by the volume provider part
1230 to the magnetic disk device 1300. The channel I/F
1200 at the host side is a block device I/F such as a
5 fibre channel I/F and the SCSI which can be connected
to the host 1100. In this embodiment, the channel of
the host 1100 side is separated from the channel of the
magnetic disk device 1300 side but the channel I/F 1210
of the host 1100 side and the channel I/F 1220 of the
10 magnetic disk device 1300 may be shared by each other.

Moreover, this volume server 1200 provides a
logical volume having a limitless capacity to the host
1100. That is, when the volume server 1200 is viewed
from the host 1100, the start value of the logical
15 block number (LBA) defining a volume range is always 0
but the LBA of the range end may not be determined.
Moreover, when the LBA of the range end is defined, the
value indicates the volume maximum capacity provided by
the volume server 1200.

20 The volume provider part 1230 in the volume
server 1200 includes an I/O detector part 1231, a
capacity controller part 1232, a segment manager part
1233, and I/O processor part 1234, and a physical and
logical address manager 1235. The segment manager part
25 1233 has a segment management table 2000 inside, and
this segment management table 2000 is used to perform
management of the storage domain of the volume provided
by the volume server 1200.

The segment management table 2000 shown in Fig. 2 includes a disk number (disk ID) 2100, a segment number (Segment Number) 2200, a starting LBA (LBA (START)) 2300, a segment size (Size) 2400, and a value of segment state (In-Use) 2500. The disk ID 2100 is a unique identification number of a disk determined by a block level protocol such as the SCSI when the magnetic disk device 1300 is connected to the volume server 1200. The segment is a minimum unit of a storage domain in each of the magnetic disk device 1300 to be managed by the volume server 1200, and the segment number 2200 is a unique identification number for management of the segment by the volume server 1200. The LBA (START) 2300 defines a physical position where the storage domain of the segment starts in the magnetic disk device 1300 of the disk ID 2100. The segment size indicates the storage domain starting at the LBA (START) 2300 in the disk. The segment state indicates whether the segment storage domain is used or not used by the volume server 1200, using two values: value 1 indicates the in-use state and value 0 indicates the not-in-use state.

Next, the physical and logical address manager 1235 performs physical and logical management between a logical volume provided by the volume server 1200 and the magnetic disk device 1300 storing data. As means for managing the logical volume and the magnetic disk device 1300 containing real data, the

physical and logical address manager 1235 has a physical and logical management table 3000 of Fig. 3. This physical and logical management table 3000 is used for management of a logical unit number (LUN) 3100, a
5 segment number (Segment Number) 3200, a LUN LBA (START) 3330, and a LUN LBA (END) 3400. The LUN 3100 is called a logical unit number indicating a unique number of a logical volume provided from the volume server 1200 to the host 1100. The segment number 3200 is a unique
10 number of a storage domain managed by the segment manager part 1233. A logical volume provided by the LUN 3100 consists of a plurality of segments and the logical block numbers (LBA) of the logical volume are successively connected in an ascending order (from
15 younger to older) in each LUN 3100 of the physical and logical management table 3000.

For example, in Fig. 3, when LUN0 (3500), the segment number 1 (3510), the segment number 2 (3520), and the segment number 3 (3530) are connected in this
20 order and are provided, as shown in Fig. 4, as a single logical volume of the LUN0 (4000) to the host 1100. The LUN LBA (START) 3300 indicates a start position of the LBA used by the segment in the LUN. The LUN LBA (END) 3400 indicates an end position of the LBA used by
25 the segment in the LUN. The I/O detector part 1231 detects LBA in each I/O request accessed from the host 1100 to the magnetic disk device 1300. The I/O processor part 1234 performs I/O processing of actual

read/write. The capacity controller part 1232 increases/decreases a capacity requested from the I/O detector part 1231 or the commander part 1132 of the host 1100.

5 In this embodiment, the volume server 1200 is arranged as an independent unit but it is also possible that the magnetic disk device 1300 has the function of the volume server 1200.

 Moreover, in this embodiment the volume
10 expansion is performed in segment basis but the volume expansion can also be performed in block basis as follows. When the segment manager part 1233 and the physical and logical address manager 1235 are considered to be, for example, a capacity-expandable
15 file system such as a Log Structured file system, the segment as the minimum unit of the storage domain in the segment management table 2000 and in the physical and logical management table 3000 can be replaced by a block which is the minimum unit of the file system
20 storage domain and the file of the file system can be considered to be a logical volume having the LUN 3100 provided from the volume server 1200. That is, even when the segment manager part 1233 and the physical and logical address manager 1235 are replaced by a file
25 system, it is possible to provide a logical volume which can be dynamically expanded in block basis.

 Next, explanation will be given on the operation of the volume server 1200, referring to Fig.

5. After an I/O request is issued from the host 1100 to the volume server 1200 (step 5001), the volume server 1200 performs processing of the received I/O request in the volume provider part 1230 (step 5002).
- 5 Upon completion of the I/O request processing, an I/O request processing completion report is issued to the host 1100 (step 5003). The host received the I/O completion report (step 5004) and the process is terminated. The volume provider part 1230 of the
- 10 volume server 1200 seems to operate in the same way as a controller of a conventional magnetic disk device 1300 but it is configured to be able to expand a volume without limit. Hereinafter, explanation will be given on the internal operation of the volume provider part
- 15 1230 having this configuration.

The volume provider part 1230 is operated by combination of the I/O detector part 1231, the I/O processor part 1234, the physical and logical address manager 1235, the capacity controller part 1232, and

20 the segment manger part 1233. In an enlarged view 1236 of the volume provider part 1230 of Fig. 1, thick lines indicate an I/O request flow and thin lines indicate a control flow. Firstly, an I/O request flowing from the host side I/F 1210 is processed by the I/O detector

25 part 1231 and then LBA of each I/O request is physically-logically converted in the I/O processor part 1234, after which read or write is performed from or to the respective magnetic disk devices 1300. This

processing flow will be detailed below.

Firstly, explanation will be given on the operation of the I/O detector part 1231, referring to Fig. 6. A currently allocated domain capacity of a logical volume in each LUN is calculated by adding the
5 respective segment size 2400 in the LUN using the segment number of the physical and logical management table 3000 of Fig. 3 and the segment management table 2000 of Fig. 2 describing the size 2400 of each segment
10 (step 6001). When an I/O request is sent from the host 1100, an LBA to be accessed in the logical volume by the I/O request is detected (step 6002). When the currently allocated area is determined to be larger than the LBA accessed (step 6003), the I/O processing is
15 requested to the physical and logical address manager 1235 (step 6005). When the currently allocated area is determined to be smaller than the LBA accessed (step 6003), then a capacity allocation is requested to the capacity controller part 1232 (step 6004) and the I/O
20 processing is requested to the I/O processor part 1234 (step 6005).

Next, referring to Fig. 7, explanation will be given on the operation of the capacity controller part 1232. The operation is different for a request
25 from the I/O detector part 1231 and for a request from the commander part 1132 of the host 1100. Firstly, explanation will be given on the operation for the I/O detector part 1231. When the I/O detector part 1231

issues a capacity increase command (step 7001), a volume increase processing is performed (step 7002) and the processing is terminated. On the other hand, when the commander part 1132 of the host 1100 issues a
5 command to reduce the entire capacity by mLBA (step 7001), the volume reduction by mLBA is performed (step 7003).

In a volume increase processing (step 7002), the segment manager part 1233 operates in combination
10 with the physical and logical address manager 1235 as follows (Fig. 8). Firstly, a free segment fetch request is sent to the segment manager part 1233 with a segment number parameter $SN = -1$ (step 8001). When the segment requested is present, the segment manager part
15 1233 puts a segment number ($SN \geq 0$) in the segment number parameter SN and returns it. When the segment requested is not found, $SN = -1$ is returned (step 8002). Upon reception of the segment number, the capacity controller part 1232 determines whether the
20 segment number parameter SN is a positive number (step 8003). If $SN \geq 0$, the process is continued and if $SN < 0$, an error is returned (step 8006), thus terminating the process. The capacity controller part 1232 makes a volume connection request to the physical and logical
25 address manager 1235 according to the segment number SN fetched from the segment manager part 1233 (step 8004). The physical and logical address manager 1235, according to the segment number, connects a segment to

the end segment of a logical volume as a target LUN in the physical and logical management table 3000 (step 8005), thus terminating the process. The operation of this physical and logical address manager 1235 will be
5 detailed later.

On the other hand, a volume reduction process (step 7003) is performed as follows (Fig. 9). When the reduction capacity is mLBA, at least one segment number having no logical volume domain after the reduction is
10 fetched from the physical and logical management table 3000 (step 9001). The segment fetch is performed as follows. Firstly, from the physical and logical management table 3000, the LBA after the reduction is calculated from the end LBA (3400) as $LBA (END) - m$.
15 According to this LBA after the reduction, one or a plurality of segment numbers (SN) not overlapping with the LBA after the reduction in the target LUN are fetched.

Next, for these segments, a segment return
20 request is made to the physical and logical address manager 1235 (step 9002). The physical and logical address manager 1235 disconnects a record matched with the segment number (SN 3200, LBA (START) 3300, LBA (END) 3400) from the list of the logical volume
25 specified by the LUN in the physical and logical management table 3000 (step 9003). The capacity controller part 1232 issues a command to manage the disconnected segment number SN as a free segment to the

segment manager part 1233 (step 9004). The segment manager part 1233 manages the disconnected segment as a free segment (step 9005). The disconnection of the logical volume is reported to the host 1100 (step 5 9006). Thus, the volume reduction process is completed.

The segment manager part 1233, upon reception of an instruction from the capacity controller part 1232, performs a segment state management of the 10 segments described in the segment management table 2000 and management of segment addition/deletion. The segment state management is performed according to the aforementioned management procedure of the segment management table 2000.

15 Referring to Fig. 10, explanation will be given on the operation of the segment state management. Firstly, a volume control request is passed from the capacity controller part 1232 together with a segment number of an argument. If the segment number < 0, it 20 is determined that the request is made for fetching a free segment (step 10001) and the process proceeds to step 10002. When the free segment fetch request is made, it is checked whether a free volume exists by using the segment management table (step 10002). If no 25 volume is free (step 10002), an error is reported to the capacity controller part 1232 (step 10005). If a free segment exists (step 10002), the segment number of the free segment is returned to the capacity controller

part 1232 (step 10003). The state for the segment in the segment management table is modified to 1 indicating the in-use state (step 10004) and the process proceeds to step 10007.

5 On the other hand, if the segment number ≥ 0 (step 10001), the request is determined to be an operation for transferring the segment not-in-use state, and the process is continued in step 10006. In step 10006, the state of the segment number is set to 0
10 indicating the not-in-use state in the segment management table 2000 and the process proceeds to step 10007. In step 10007, by using items of the state 2400 in the segment management table 2000, a segment use ratio (the number of segments in use against the total
15 number of segments) is calculated from the number of segments in use and the total number of segments and it is determined whether the use ratio is equal to or more than 90 %. If the use ratio is 90 % or above, a request is made to a maintenance staff to add a
20 magnetic disk device 1300 (step 10008), thus terminating the process. If the use ratio is below 90%, then the process is terminated as it is. The threshold value of the use ratio may also be set to other than 90% according to the system reliability by
25 the maintenance staff. This completes the explanation on the operation of the volume use state management by the segment manager part 1233.

Next, referring to Fig. 11, explanation will

be given on the addition/deletion operation of the magnetic disk device 1300. Step 11001 determines whether a volume addition or deletion is to be performed. If addition, then the process proceeds to step 11005; and if deletion, the process proceeds to step 11002. In case of addition of the magnetic disk drive 1300, step 11005 checks the capacity of the magnetic disc device 1300 (step 11005). According to the capacity checked, the number of segments and the size are determined (step 11006). The number of segments and the size may be values specified by a maintenance staff or fixed values. According to the values, upon each addition of a magnetic disc drive 1300 to the end of the segment management table 2000, the disk ID 2100, the segment number 2200, the start position 2300, the segment size 2400, and 0 (not-in-use) as the use state 2500 are inserted (step 11007), thus completing the process.

On the other hand, in case of a volume deletion, it is determined that the segment of the segment number is a free segment or not (step 11002). If the segment is free, a record of that segment number is deleted from the segment management table 2000 (step 11003). If the segment is not a free segment (step 11002), an error is returned and the process is terminated (step 11004).

This completes the explanation on the segment addition/deletion operation of the segment manager part

1233. The segment manager part 1233 performs management of the segment use state and the addition/deletion operation according to the aforementioned procedure.

5 Referring to Fig. 12, explanation will be given on the operation of the I/O processor part 1234. Firstly, an I/O request sent from the I/O detector part 1231 to the I/O processor part 1234 requests the physical and logical address manager 1235 to convert
10 the LBA of the I/O request into a physical address (step 12001). Next, a type of the I/O request is determined (step 12002). If the I/O request is a read I/O, using the physical address, data is read out from the magnetic disk device 1300 (step 12003) and the read
15 out data is returned to the host 1100 (step 12004). If the I/O request is a write I/O, using the physical address, data is written to the magnetic disk device 1300 (step 12005) and the write completion is reported to the host 1100 (step 12006). This completes the
20 explanation on the operation procedure of the I/O processor part 1234.

 The operation of the physical and logical address manager 1235 can be separated in two processes. One of them is a physical and logical address
25 conversion procedure requested by the I/O processor part 1234. This conversion procedure will be explained with reference to Fig. 13 and then a physical-logical structure between an actual physical address and a

logical address will be explained with reference to Fig. 14. For the conversion, a rule as follows is used. Firstly, a segment group corresponding to the LUN as an I/O request target is selected (step 13001).

- 5 A disk number and an LBA of the physical side are calculated from the LBA (step 13002).

The calculation is performed as follows.

- Firstly, according to an LBA target described in the I/O request, by referencing each segment start LBA from
10 the physical and logical management table 3000, a segment number (SN) to be accessed by the LBA target (143000 in Fig. 14) is selected from a plurality of segments to which the target LUN belongs. According to the selected segment number SN, using the segment
15 management table 2000, a disk ID to be accessed by the LBA target (14300), and a physical address of the segment from the head of the magnetic disk device are specified. When viewed from the N-th segment head physical address LBAN (14003), the address on the
20 segment where data is operated is LBA target - LBAN. Accordingly, the physical address of the segment from the head of the magnetic disk device 1300 is: (segment head LBA LBA segstart) + (LBA target - LBAN) (14005). The disk ID 2100 and the value of (segment head LBA LBA
25 segstart) + (LBA target - LBAN) (14005) as the LBA of the magnetic disk device 1300 of the disk ID 2100 are returned to the I/O processor part 1234 (step 13003).

The other operation process of the physical

and logical address manager 1235 is volume addition/deletion to/from the physical and logical management table 3000. This operation will be explained with reference to Fig. 15. A LUN in the
5 physical and logical management table is selected (step 15001). For a volume addition (step 15002), a segment number to be added is given from the capacity controller part 1232 and in the LUN volume to be added, the segment is added to the end of the segment record
10 (step 15003). Here, when the added segment is viewed from the host 1100, an added volume exists between the addresses from LBAN (14003) to LBAN + SIZEN (14004). When deleting a volume (step 15002), after receiving an instruction of deleting mLBA from the capacity
15 controller part 1232, the capacity of mLBA is deleted in segment basis from the end of the segment record in the LUN volume to be deleted (step 15004). The segment deleted returns its segment number to the capacity controller part 1232 (step 15005). This completes the
20 explanation of the operation of the physical and logical address manager 1235.

The commander part 1132 of the host 1100 is a program which, after reducing the data capacity in the volume used by the application 1110, upon reception of
25 an instruction from the application or the user, sends an instruction to the volume server 1200 via the I/F 1140 so as to reduce in segment basis the storage domain of the logical volume provided by the volume

server 1200. This operation will be detailed below by referring to Fig. 16.

Firstly, a request to reduce the volume by mLBA is received from the application 1110 (step 5 16001). The commander part 1132 reduces by mLBA the end block position of the logical area information of the logical volume under control of OS 1130 (step 16002). The volume of mLBA capacity is issued from the commander part 1132 to the capacity controller part 10 1232 of the volume server (step 16003). The capacity controller part 1232 performs a volume reduction processing (step 16004) and returns the result to the host 1100. This completes the explanation of the operation of the commander part 1132.

15 Next, explanation will be given on the interlocked operation with the application 1110. As has been described above, the host 1100 can handle a logical volume provided from the volume server 1200 as an existing volume. However, when a volume exceeding 20 the segment held by the volume indicated by the LUN is accessed in the volume server 1200, a volume is added from the segment manager part 1233. This operation is identical to the operation of the aforementioned volume provider part 1230.

25 On the other hand, when the application 1110 has performed a volume capacity control of a volume expanded by itself and the capacity actually used by the application 1110 has reduced, a capacity difference

before and after the deletion is reported to the commander part 1132 of the host 1100 by the application 1110 and the aforementioned volume reduction procedure (Fig. 16) is performed in the volume server via the
5 commander part 1132 of the host 1100. By these processes, the number of segments in use by the volume server is matched with the operation of the application 1110 and an appropriate logical volume capacity can be provided.

10 (2) Second Embodiment

In the first embodiment, when a write I/O access occurs to other than a storage domain of a logical volume provided from the volume server 1300 to the host 1100, the volume server 1200 successively
15 allocates a storage domain, thereby performing a storage domain expansion to the logical volume provided by the volume server 1200. In the second embodiment, according to an instruction of the application 1110, a storage capacity expansion of a logical volume is
20 performed in advance via the commander part 1132 of the host 1100 and the storage capacity of the logical volume is recognized by the OS 1130 before becoming usable by the application 1110. This procedure will be explained by referring to Fig. 18.

25 Firstly, the commander part 1132 of the host 1100 receives a storage domain expansion request by mLBA from the application 1110 (step 18001). The commander part 1132 of the host 1100 issues a write I/O

to the LBA as a total of the entire logical volume capacity and the expanded domain mLBA (step 18002). The I/O detector part 1231 of the volume server 1200 performs a capacity expansion processing to the
5 capacity controller part 1232 because the volume provided by the volume server is expanded by mLBA (step 18003). If the write I/O is successful (step 18004) and the storage domain is increased, then step 18005 is performed, and if the write I/O fails, the process
10 terminates without performing the storage domain expansion. In step 18005, the commander part 1132 of the host 1100 increases by mLBA the end block count of the logical area information of the OS 1130 corresponding to the logical volume provided by the
15 volume server 1200. This completes the explanation of the storage domain expansion procedure of the logical volume provided from the volume server 1200 via the application 1110.

According to the present invention, by adding
20 a storage domain to a volume server according to a request from the volume server, the application of the host computer can use a single volume which can be dynamically expanded.